# Monitoring and Maintaining Exchange 2000 Server

Monitoring system activity and server performance, when combined with disaster planning and regular backup, is a necessary part of preventive maintenance of your Microsoft® Exchange 2000 Server deployment. Through monitoring, you obtain data that you can use to diagnose system problems, plan growth, and troubleshoot problems. Using Exchange 2000's Monitoring and Status tool, diagnostic and extended logging, and Queue Viewer, along with Microsoft® Windows 2000-based tools such as Performance Monitor, Event Viewer, Task Manager, and Terminal Services Client, ensures that you have up-to-date information about how your Exchange Server and network are operating. Two additional services, Network Monitor and DOS-based netdiag, provide additional network monitoring information.

> **Note** Because of architectural difference between the monitoring interfaces in Exchange 5.5 and Exchange 2000, mixed mode networks require that you use both monitoring systems. Use the Exchange 2000 monitoring system to monitor Exchange 2000 only, and use the Exchange 5.5 monitoring system to monitor Exchange 5.5 only. You can monitor using only the Exchange 2000 monitoring system, when you are in native mode.

**In This Chapter**

Exchange 2000 Monitoring Features

Windows 2000 Performance Monitoring Tools

Analyzing Performance Data

# Exchange 2000 Monitoring Features

Exchange 2000 provides a number of features that assist you in monitoring and maintaining your Exchange server and network. These tools include the Monitoring and Status Tool, Diagnostic Logging and Viewing Tool, and Extended Logging Tool.

## Monitoring and Status Tool

The Monitoring and Status tool, located in the tools folder in the Exchange System Manager, is the primary Exchange 2000 tool you use to keep track of the health and status of your network and servers. The tool is composed of two interfaces: Notifications and Status.

### Notifications

You can use the Notifications interface in conjunction with the Status interface to set up email alerts or script triggers when a warning or critical state is reached on any of your network servers. In the Status interface, you configure warning and critical states for an array of Microsoft® Windows NT 4.0, Windows 2000, and Exchange 2000 services, SMTP and X.400 Queues, Available virtual memory, CPU utilization, and free hard disk drive space.

### Configuring Warning and Critical States

You can configure Exchange to constantly monitor the performance levels of an array of network and application services. Levels for both warning states and critical states can be established so that problems or potential problems are announced and can be dealt with as they arise. You configure these through the Status and Notifications user interface.

**Note** You can also configure Exchange to monitor performance levels using the server properties Monitoring tab of the server properties menu in the Status tool in Exchange System Manager.

Effective monitoring requires that you establish levels of acceptable performance for each resource. From this baseline level you can determine a warning threshold and critical state threshold.

-

Exchange 2000 logs a critical state warning when any of the following Default Exchange Services stop running:

- Information Store
- MTA Stacks
- Routing Engine
- System Attendant
- SMTP
- World Wide Web Publishing Service

You can add other services to this set from the Add Service menu of the **Details** pane under the **Default Microsoft Exchange Services** selection. Any service that you add to the **Default Microsoft Exchange Services** follows the same configuration rules as the set. This means that a critical error arises if the selected service stops running. If you want to add a set of services so that a warning message rather than a critical error message logs, then you must create another set. To do this, select **Add** in the **Add Resource** pane on the **Monitoring** tab, choose Windows 2000 Service, then select the desired service or services from the **Details** pane.

For more information about the Windows 2000 Services that you can monitor, see the "Windows 2000 Services Available For Monitoring" appendix in this book.

**Note** You use most of these services exclusively for troubleshooting and they do not require monitoring on a properly functioning server. Many services start and stop frequently in normal operating conditions; therefore, use consideration before configuring Exchange 2000 to generate errors when a service stops running. Also, because monitoring and generating notifications consumes server resources, it is recommended that you do not configure Exchange 2000 to monitor an excessive number of services.

### Configuring Notifications

Because you cannot instantly access all elements of every network server, Exchange provides a Notifications tool that you can set to trigger a script or alert appropriate personnel when Exchange crosses a warning or critical threshold.

**Note** Exchange 2000 is not configured by default to fire any notifications; all notifications must be configured through the User Interface.

For example, you can set a 5-minute warning threshold and a ten-minute critical threshold for SMTP queue growth and configure Exchange to notify you through e-mail. When the SMTP queue growth reaches 5 minutes, Exchange sends the following message by default:

%TargetInstance.Name% has reported a %TargetInstance.ServerStateString%. Reported status is:

Queues - %*TargetInstance.QueuesStateString*%

Drives - %*TargetInstance.DisksStateString*%

Services - %*TargetInstance.ServicesStateString*%

**Note** %*TargetInstance.Name*%, %*TargetInstance, ServerStateString*%, and other fields are placeholders into which Exchange places data about the specifics of the warning, including the server name and instance data.

For more information about setting Notifications, see Exchange 2000 Help.

## Status

The Status interface allows you to view servers and connectors on your network with their status condition and Admin Group designation. You can also disable all monitoring of a server through this interface.

Monitoring the status of servers and connectors is one of the most effective ways to ensure that your system and network functions properly. The status pane of the Exchange System Monitoring and Status Tool shows the operating status of each server and connector on the network.

The following status states and definitions apply to servers:

- Available: Server on-line and functioning normally
- Unreachable: One of the main services on the server is down

**Note** If a server is unreachable and is in a different Routing Group, it may indicate that a connector between Routing Groups is down or does not exist.

- In Maintenance Mode: Monitoring is disabled on this server for maintenance, backup, repair, or other reason.
- Unknown: System Attendant cannot communicate with the local server

The following status states and definitions apply to connectors:

- Available: Functioning properly
- Unavailable: A communication function such as routing service is not functioning on this connector

When you want to disable monitoring on a server for maintenance, backup, or recovery, select the **Disable all monitoring on this server** checkbox on the **Monitoring** tab of the server properties menu for an individual server selected from the status pane in the Exchange System Manager Status tool interface.

# Diagnostic Logging and Viewing Tool

Using the Diagnostic Logging feature in Exchange 2000 to monitor protocol connectors and Microsoft Exchange connectors is an effective way to keep up-to-date on a server's status and potential trouble. Diagnostic Logging produces information on monitored connectors that you can view using Event Viewer, which is located in the Windows 2000 Administrative Tools.

Exchange allows you to log diagnostic data using The Services and Categories windows on the Diagnostics Logging tab of the Properties menu for each physical server. You gain access to this menu through Exchange System Manager. Table X.X provides a listing of the services available for monitoring and their abbreviation as shown in the Services window.

| Service To Monitor | Abbreviation to Select |
|---|---|
| IMAP4 Protocol | IMAP4Svc |
| Lotus Notes Group Wise | LME-GWISE |
| Lotus Notes | LME-NOTES |
| Address List | MSExchangeAL |
| Lotus CC Mail | MSExchangeCCMC |
| Microsoft Exchange Directory Synchronization | MSExchangeDX |
| Microsoft Exchange SchedulePlus Free-Busy | MSExchangeFB |
| Microsoft Exchange Router for Novell GroupWise | MSExchangeGWRtr |
| Information Store System | MSExchangeIS\System |
| Information Store Mailbox | MSExchangeIS\Private |
| Information Store Public Folders | MSExchangeIS\Public |
| Microsoft Mail | MSExchangeMSMI |
| MAPI Address Book Proxy Service | MSExchange NSPI Proxy |
| MAPI Address Book Referral Service | MSExchangeRFR Interface |
| Site Replication Service | MSExchangeSRS |
| SMTP Routing Engine and Transport | MSExchangeTransport |
| POP3 Protocol | POP3Svc |

Table X.X Services Available for Diagnostic Logging

Each Service contains at least one category in which you can establish a logging level. This is similar to the Performance Monitor Performance Objects and Counters interface. For instance, the POP3Svc Service contains the following categories: Connections, Authentication, Client Action, Configuration, Content Engine, and General. You can

select different logging levels for each of the service categories, which allows you to derive precise system performance information.

You can set logging levels for most services and categories.  When a connector tries to print a log message, it first checks the message's logging level against the logging level you configure.  If the message's logging level is the same or more severe than the configured one, the message is logged in Event Viewer.  Otherwise, the connector does not log a message and continues.

Logging level options are as follows:

None: only error messages are logged

Minimum: warning messages and error messages

Medium: informational messages, error messages, and warning messages

Maximum: trouble-shooting messages (fine and detailed info), error messages, warning messages, and informational messages

**Note**  It is not recommended that you use the Maximum logging setting unless instructed to do so by Exchange technical support because logging at Maximum considerably drains resources.

# Protocol Logging Tool

The Protocol Logging Tool provides detailed information about the commands being sent and received by SMTP and NNTP.  This tool is particularly useful in monitoring and troubleshooting protocol or messaging errors.

The User Interface for SMTP and NNTP protocol logging is located in the **Properties** menu of an individual SMTP or NNTP virtual server. To configure SMTP and NNTP logging, select the desired virtual server from the **Protocols** folder, right-click and select **Properties**. The General tab in Extended Logging Properties is shown in Figure X.X.

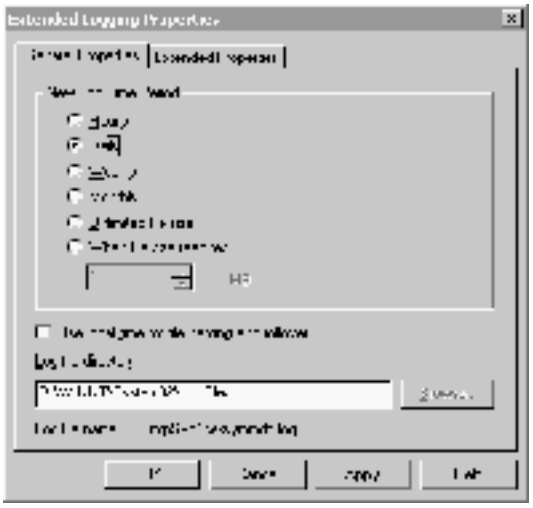

Figure X.X: General tab in Extended Logging Properties

When you click the **Properties** button on the **General** tab of the **Virtual Server Properties** dialog box, you can select the logging frequency as well as the name and location of the log file to create. By default, the log file path is as follows:

system Root:\WINNT\System32\LogFiles\

You can change the path or log file name by clicking one of the option buttons or by letting the current path stand as the default setting. You must select a logging format for the SMTP or NNTP log. By default, the log file name uses the following creation date For SMTP server:

SmtpSvc1\exyymmdd.log. Figure X.X shows the Extended Properties tab for Extended Logging Properties.

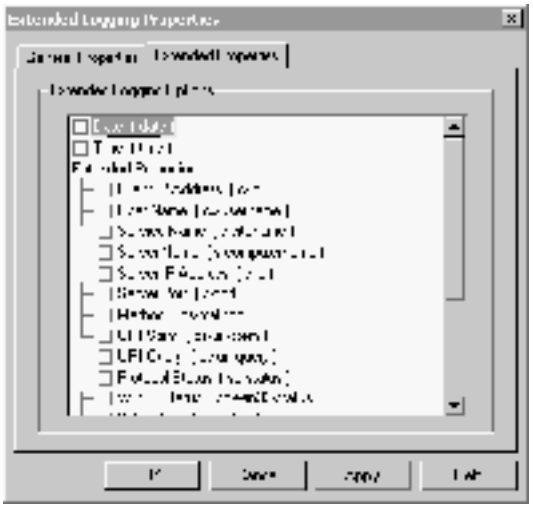

Figure X.X: Extended Properties tab in Extended Logging Properties

On the **Extended Properties** tab, you can select the configuration parameters under the **Extended Logging Options** menu. These logging options provide more detailed information on the Service and Category logging properties already configured for the virtual server. They do not add additional services or categories to log. Unlike the IMAP4, POP3, and MSExchange connectors, you cannot select the services and categories to log for SMTP or NNTP.

You can also establish a rollover timeframe or file size under the New Log Time Period section of  the General Properties page of the Extended Logging Properties Tab. Logs will be overwritten when the time interval expires or when the log file reaches the set size.

## Using Event Viewer to View Logs

Event Viewer is an MMC snap-in that provides event information about running Applications, Directory Service, File Replication Service, Security, and the System. It also allows you to view the logs you configure for IMAP4, POP3 protocols and MSExchange connectors. Events are logged by date, time, source, category, event number, user, and computer. By viewing the event data, you can view errors, warnings, and diagnostic information to find problems that occur on any computer in the network.

By viewing an event's property, you can see the event's logging properties and text. For more information about using Event Viewer, see Windows 2000 online Help.

## Queue Viewer

As part of the monitoring process, you can view the X.400, SMTP transport protocol queues and other connectors installed on a server by using Exchange System Manager. You can use information, such as message age, the number of messages in queue, and other information to troubleshoot problems on a server.

You access the queues through the paths, as shown in Figure X.X

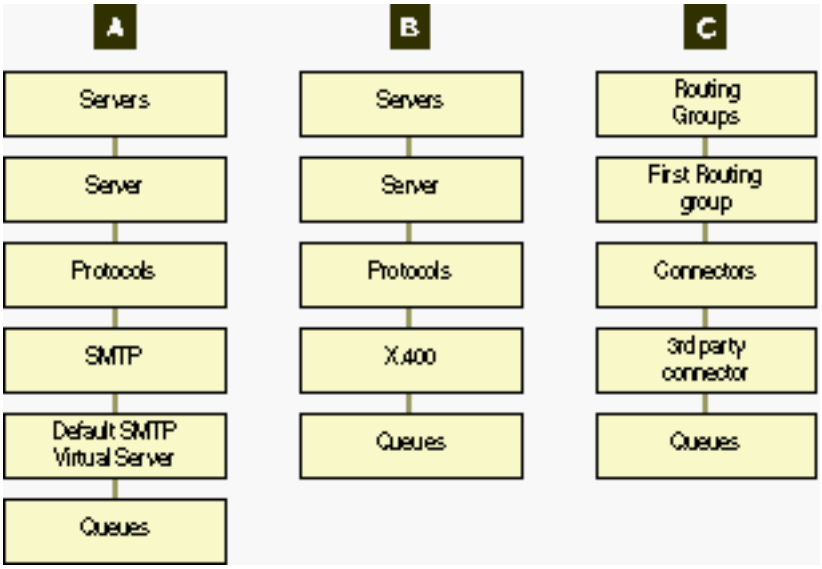


Figure X.X: Paths to Accessing Queues

The two queues most useful to monitor are the Local Delivery and Messages Awaiting Directory Lookup. A backlog in the Local Delivery Queue indicates trouble with the Store. A backlog in the Messages Awaiting Directory Lookup Queue indicates that there is a problem contacting the domain controller.

Additionally, a backlog of a number of messages with the same destination indicates there may be a problem with the destination domain controller.

# Windows 2000 Performance Monitoring Tools

Exchange 2000 is integrally linked to Windows 2000, so that many of the tools you use to monitor your Exchange Server and network are part of the Windows 2000 operating system. Performance Monitor, Netdiag, Task Manager, and Network Monitor are all Windows 2000 tools that you use to monitor performance.

## Performance Monitor

System performance in Exchange is monitored in part by the Performance Console Module (perfmon), which includes System Monitor and Performance Logs and Alerts.

## System Monitor

The System Monitor displays configurable graph, histogram, or report displays. System Monitor provides short-term viewing of data, troubleshooting, and diagnosis. It includes tools such as physical hard disk drive counters and workload balance tools.

Performance objects, which are logical counters that are associated with a resource or service that you can monitor, contain at least one performance counter. You can view selected performance counters individually or in relation to other available counters.

**Note** Monitoring large numbers of counters can create a high amount of overhead, which can make the system unresponsive to keyboard or mouse input. To reduce this burden, display data in report view when collecting information from large numbers of counters, or direct data to a binary log and view the data in system Monitor as it is written to the log.

For more information about System Monitor features, see Windows 2000 Help

## Performance Objects

Much of the performance monitoring you do in Exchange 2000 relies on data from Windows 2000 performance objects, which obtain performance data from components in your computer and are monitored by Windows 2000 System Monitor and Performance Logs and Alerts. As a system component functions in your system, it generates performance data. The data is formulated into performance objects and is typically named for the component generating the data. For example, the Processor object is a collection of performance data about processors on your system.

A range of performance objects are built into the Windows 2000 operating system, typically corresponding to the major hardware components such as memory, processors, and so on; other programs might install their own performance objects.

Exchange 2000 installs its own set of performance objects and counters. The Exchange 2000 performance objects subset listed in table X.X provide information about the services or resources in Exchange that you can monitor using System Monitor in the Performance Monitor MMC.

| Feature or Service to Monitor | Performance Object Available |
|---|---|
| Active Directory DXA Connector | MSExchangeADDXA |
| Address List | MSExchangeAL |
| Chat Communities | MSExchange Chat Communities |
| Chat Service | MSExchange Chat Service |
| Directory Service Access Cache | MSExchangeDSAccess Caches |
| Directory Service Access Contexts | MSExchangeDSAccess Contexts |
| Directory Service Access Processes | MSExchangeDSAccess Processes |
| Document Conferences | MSExchangeCONF |
| Document Conferencing Manager | MSExchangeDcsMgr |
| Document Conferencing Protocol | MSExchangeT.120 |

| | |
|---|---|
| (Multipoint Control Unit) | |
| Epoxy Queues and Activity | EXIPC |
| Event Store | MSExchangeES |
| File Replication Connector | FileReplicaConn |
| File Replication Settings | FileRepSet |
| HTTP Extension | Exchange Server HTTP Extension |
| Internet Information Server Store Driver | Exchange Store Driver (IIS) |
| Internet Message Access Protocol Version 4 | MSExchangeIMAP4 |
| Information Store | MSExchangeIS |
| Private Information Store | MSExchangeIS Mailbox |
| Public Information Store | MSExchangeIS Public |
| System Information Store | MSExchangeIS |
| Mailbox Information Store | MSExchangeIS Mailbox |
| Public Folders Information Store | MSExchangeIS Public |
| Lotus CC Mail | MSExchangeCCMC |
| Lotus Notes Message Center | MSExchangeNMC |
| Message Transfer Agent | MSExchangeMTA |
| Message Transfer Agent Connections | MSExchangeMTA Connections |
| MS Mail Connector Interchange | MSExchangeMSMI |
| Exchange Referral Service | MSExchangeSA-RFR |
| MS Mail Connector Mail Transfer Agent | MSExchangePCMTA |
| Name Service Provider Interface (Active Directory Integration) | MSExchangeSA-NSPI Proxy |
| Network News Transfer Protocol Commands | NNTP Commands |
| Network News Transfer Protocol Server | NNTP Server |
| Novell Groupwise Connector | MSExchangeGWC |
| Object Linking and Embedding database events | MSExchangeOledb Events |
| Object Linking and Embedding database resources | MSExchangeOledb Resources |
| Post Office Protocol Version 3 | MSExchangePOP3 |
| Service Account | MSExchangeSA |
| Site Replication Service | MSExchangeSRS |

| | |
|---|---|
| Simple Mail Transfer Protocol | SMTP |
| Store Driver | Exchange Store Driver (Store) |
| Video Conferencing | MSExchangeIPConf |
| Web Mail | MSExchangeWebMail |

Table X.X Exchange Services or Resources Monitored Using System Monitor

Each Exchange 2000 Performance Object has at least one associated counter, which you can configure to monitor instances involving any number of server actions. For more information about the function of a counter for a specified performance object, select the **Select Counters from List** button, highlight the desired counter in the **Select Counters from List** window and click **Explain**.

## Physical Hard Disk Drive Counters

Hard disk drive usage statistics help you balance the workload of network servers. System Monitor provides physical hard disk drive counters for troubleshooting, capacity planning, and measuring activity on a physical hard disk drive.

At a minimum you should monitor the following counters:

- Physical Disk\ Disk Reads/sec and Disk Writes/sec
- Physical Disk\ Current Disk Queue Length
- Physical Disk\ % Disk Time
- LogicalDisk\ % Free Space

When testing hard disk drive performance, log performance data to another hard disk drive or computer so that it does not interfere with the hard disk drive you are testing.

You might want to observe additional counters, include the following counters:

- Physical Disk\ Avg. Disk sec/Transfer
- Physical Disk\ Avg. Disk Bytes/Transfer,
- Physical Disk\ Disk Bytes/sec.

The Physical Disk\ Avg. Disk sec/Transfer counter reflects how much time a hard disk drive takes to fulfill requests. A high value might indicate that the hard disk drive controller is continually retrying the hard disk drive because of failures. These failures increase average hard disk drive transfer time. For most hard disk drives, high average hard disk drive transfer times correspond to values greater than 0.3 seconds.

You can also check the value of Avg. Disk Bytes/Transfer. A value greater than 20 KB indicates that the hard disk drive is performing well. Low values result when an application accesses a hard disk drive inefficiently. For example, applications that access a hard disk drive at random raise Physical Disk\ Avg. Disk sec/Transfer times because random transfers require increased seek time. Disk Physical Disk\ Bytes/sec provides you with the throughput rate of your hard disk drive system.

Because hard disk drive counters can cause a modest increase in hard disk drive access time, Windows 2000 does not automatically activate the counters at system startup.

## Logical Disk Counters

Unlike physical hard disk drive counter data, logical hard disk drive counter data is not collected by the operating system by default. To obtain performance counter data for logical drives or storage volumes, you must type **diskperf -yv** at the command prompt. This causes the hard disk drive performance statistics driver used for collecting hard disk drive performance data to report data for logical drives or storage volumes. By default, Windows 2000 uses the **diskperf -yd** command to obtain physical drive data. For more information about using the **diskperf** command, type **diskperf -?** at the command prompt.

## Determining Workload Balance

To balance loads on network servers, you need to know how busy the server hard disk drives are. Use the Physical Disk\ % Disk Time counter, which indicates the percentage of time a drive is active. If the results of Physical Disk\  % Disk Time is high (over 90 percent), check the Physical Disk\ Current Disk Queue Length counter to see how many system requests are waiting for hard disk drive access. The waiting Input/Output (I/O) requests should be sustained at a number no more than 1.5 to 2 times the number of physical hard disk drive spindles.

Most hard disk drives have one spindle, although Redundant Array of Inexpensive Disks (RAID) devices usually have more. A RAID device appears as one physical hard disk drive in System Monitor. RAID devices created through software appear as multiple drives (or instances). You can either monitor the Physical Hard disk drive counters for each physical drive (other than RAID), or you can select **All Instance** to monitor data for all the computer's drives.

Use the values of the Current Disk Queue Length and % Disk Time counters to detect bottlenecks in the hard disk drive subsystem. If Current Disk Queue Length and % Disk Time values are consistently high, consider upgrading the hard disk drive or moving some files to an additional hard disk drive or server.

The system maps physical drives to logical drives using the same instance name. Therefore, if you have a dynamic volume that consists of multiple physical hard disk drives, instances might appear as **Disk 0 C:**, **Disk 1 C:**, and **Disk 2 D:**, (where C: is made up of physical drives 0 and 1). If you have two logical partitions on a hard disk drive, the instance appears as **0 C: D:**.

For hardware-enabled stripe sets, statistics for each hard disk drive are not available. You can obtain this data only when monitoring stripe sets enabled in software.

If you use a RAID device, the % Disk Time Counter can indicate a value greater than 100 percent. If this happens, use the Avg. Disk Queue Length counter to determine the average number of system requests waiting for hard disk drive access.

# Performance Logs and Alerts

Performance Logs and Alerts contain features for logging counter and event tracing data and for generating performance alerts. With counter logs, you can record data about hardware usage and the activity of system services from local or remote computers. You can configure logging to occur manually on demand, or automatically based on a user-defined schedule. You can enable continuous logging, which is subject to file size

or duration limits. You can view logged data using the System Monitor display or you can export it to a spreadsheet program or database for analysis and to generate a report.

Trace logs record data when an event such as a hard disk drive Input/Output (I/O) error or page fault occurs. When the event occurs, the provider sends the data to the log service.

You can set an alert on a counter, which can send a message, run a program, or start a log when the selected counter's value equals, exceeds, or falls below a specified setting.

For more information about Performance Logs and Alerts features, see Windows 2000 help.

To see the status of a tool, select it an If any logs or alerts are defined, then they appear in the **Details** pane. A sample settings file for a counter log named System Overview is included with Windows 2000 and appears in the **System Overview** pane. You can use this file to see basic system data such as memory, hard disk drive, and processor activity.

Right-click the **Details** pane to create a new log or alert. You can do this in a new file or you can use settings from an existing HTML file as a template.

If you want to configure a counter log or an alert, use the **Add Counters** dialog box to specify objects, counters, instances, and updates. If you want to configure a trace log, use the **General** property tab. Each tool offers unique properties, but some options might not be available for all tools.

# Task Manager

Task Manager provides information about applications running on your system, processes, memory usage and statistics about memory and processor performance. However, it lacks the logging and alert capabilities of the Performance console and does not provide the breadth of information available from System Monitor counters.

## Monitoring Processes

To use Task manager to monitor a process, click the Task Manager's **Processes** tab to see a list of processes that are running and data about their performance. Task Manager process tables include all processes that run in their own address space, including all applications and system services. Exchange 2000 runs two processes that you can view using Task Manager: store.exe and inetinfo.exe.

> **Note** Even if you have more than one virtual server running on a server, only one instance of store.exe and inetinfo.exe exists.

## Monitoring the System

To use Task Manager to see a dynamic overview of system performance, including a graph and numeric display of processor and memory usage, click the Task Manager Performance tab.

To graph the percentage of processor time in privileged or kernel mode, click the **View** menu's **Show Kernel Times** button. This is a measure of the time that applications to

operate system services. The remaining time, known as user mode, is spent running threads within the application code.

If you use multiple-processor computers, you can click the **View** menu's **CPU History** button and view a graph of the non-idle time of each processor in a single graph or in separate graphs.

For more information about Task Manager, see Windows 2000 help.

# Terminal Services Client

The Terminal Services Client interface, located in Windows 2000 Administrative Tools, allows you to log on to a remote network computer from another terminal. You can use this interface to do remote administration and monitoring of any accessible server on the network.

For example, if you are in Seattle and you discover that a server in Tokyo is not responding, and you can't contact the on-site administrator, but the server has Terminal Services Client installed, you can remotely log onto the computer and administer it. You can restart it, run installed applications, start and stop services, execute monitoring procedures, and so on.

Another benefit of Terminal Services Client is that you can run applications that are not installed on your computer by gaining access to a computer with Terminal Services Client installed that has the application installed.

For more information on Terminal Services Client, see Windows 2000 Support Tools Help.

# Network Monitor

Monitoring a network typically involves observing resource use on a server and measuring overall network traffic. You can use Network Monitor to understand the traffic and behavior of your network components. Unlike System Monitor, which you use to monitor hardware and software, Network Monitor exclusively monitors network activity. You can use System Monitor to perform both of these functions, however for in-depth traffic analysis, you should use Network Monitor.

## Observing Resource Use

To check resource use, you should start by tracking the counters to observe resource usage on your server. To focus on network related resource usage, add the counters that correspond to the various layers of your network configuration. Abnormal network counter values often indicate problems with a server's memory, processor, or hard disk drives. For this reason, the best approach to monitoring a server is to watch network counters in conjunction with the following Performance Monitor counters for the Processor: % Processor Time, PhysicalDisk\ % Disk Time, and Memory\ Pages/sec. For more information about these counters, see "Physical Hard Disk Counters" earlier in this chapter.

For example, if a dramatic increase in Pages per second is accompanied by a decrease in total Bytes per second handled by a server, the computer probably lacks physical memory for network operations. Most network resources, including network adapters and protocol software, use nonpaged memory. A computer can page excessively if most

of its physical memory is allocated to network activities, which leaves a small amount of memory for processes that use paged memory. To verify this, check the computer's system event log for entries indicating that it is out of paged or nonpaged memory.

## Measuring Network Traffic

You can use Network Monitor to observe throughput across network layers. Investigating network performance includes monitoring activity at different network layers. There are four layers in which you monitor network activity: the data-link layer, the network layer, the transport layer, and the presentation/program layer.

### Data-Link Layer

The data-link layer includes the network adapter. Use the following network interface object counters to monitor network activity at the data-link layer:

- Bytes total/sec
- Bytes sent/sec
- Bytes received/sec.

### Network Layer

For the Network layer, use the IP object counters to monitor network activity at the network layer:

- Datagrams Forwarded/sec
- Datagrams Received/sec
- Datagrams/sec
- Datagrams Sent/sec.

### Transport Layer

The transport layer varies with network protocol in use. For TCP/IP, use the TCP object counters to monitor network activity at the transport layer:

- Segments Received/sec
- Segments Retransmitted/sec
- Segments/sec
- Segments Sent/sec.

If the retransmission rate is high, you might have a problem with your hardware.

The ICMP and UDP object counters are useful for more extensive monitoring of TCP/IP network transmissions. The ICMP performance object consists of counters that measure the rates at which Internet Control Message Protocol (ICMP) messages are sent and received by using the ICMP protocol. It also includes counters that monitor ICMP protocol errors. The UDP performance object consists of counters that measure the rates at which User Data Protocol (UDP) datagrams are sent and received using the UDP. It includes counters that monitor UDP errors.

If you use the NetBEUI protocol, use the following counters:

- NetBEUI\ Frame Bytes Received/sec
- NetBEUI\ Frames Received/sec
- NetBEUI\ Frames Rejected/sec
- NetBEUI Resource\ Times Exhausted

If you use the NWLink protocol, three objects are available: NWLink IPX and NWLink NetBIOS (for computers that communicate using the IPX protocol) and NWLink SPX (for computers that connect using the SPX protocol). Note that frame-related counters for these objects report only zeroes.

**Presentation/Program Layer**

For the presentation/program layer, use the Server object counters if you monitor a server, or the Redirector object counters if you monitor a client computer. Exchange 2000, as well as some program-layer processes, such as Web servers, might have their own object counters that you use to monitor transmissions across this layer.

The Redirector object counters collect data about requests transmitted by the Workstation service. The Server object counters collect data about requests received and interpreted by the Server service.

At a minimum, you should include the Total Bytes per second counter for both the Redirector object (for client computers that you monitor) and the Server object (for server computers).

Each of these objects provides several other counters that you might want to monitor if you suspect problems with either the Workstation or Server services:

- Redirector\Current Commands
- Redirector\Network Errors/sec
- Redirector\Reads Denied/sec
- Redirector\Writes Denied/sec
- Redirector\Server Sessions Hung
- Server\Sessions Errored Out
- Server\Work Item Shortages
- Server\Pool Paged Peak
- Server\Nonpaged Pool Failures

The Sessions Errored Out counter reports automatic disconnections and sessions that end because of an error. For more accurate values for sessions that end because of an error, obtain the value for Sessions Timed Out and reduce the Sessions Errored Out value by that amount.

# Network Diagnosis Tool (netdiag)

You can use the Windows 2000 support tool, Netdiag to diagnose problems on a network. This command-line diagnostic tool helps isolate networking and connectivity problems by performing a series of tests to determine the state of your network client.

These tests, and the network status information they provide, give you a more direct means of identifying and isolating network problems. Because this tool does not require you to specify parameters or switches, you can focus on analyzing the output, rather than training users how to use the tool. For more information about installing or using netdiag, see Windows 2000 Support Tools Help.

# Analyzing Performance Data

In general, performance monitoring concentrates on how the operating system and any applications or services use the resources of the system, such as the hard disk drives, memory, processors, and network components.

The data you accumulate through daily monitoring provides you with the information you need to analyze trends and plan your system capacity. Even if your system operates satisfactorily today, it is important to plan for changes in demand by new users, or by technologies and programs you deploy. Unanticipated network growth can result in overused resources and poor levels of network service. By characterizing system performance over time, you can justify the need for new resources before the need becomes critical. The following are some definitions to help you understand performance-monitoring terminology.

***Throughput*** is a measure of the work done in a unit of time, typically evaluated from the server side in a client/server environment. Throughput tends to increase as the load increases up to a peak level. It then begins to fall, and a queue might develop. Throughput in an end-to-end system, such as client/server, depends on how each component performs. The slowest point in the system sets the throughput rate for the system as a whole. Often this slow point is referred to as a bottleneck. Performance monitoring tells you where bottlenecks occur in your system. The resource that shows the highest use is often the bottleneck, but not always. It can also be a resource that successfully handles a lot of activity. As long as no queues develop, there is no bottleneck.

***Queue*** is a group of jobs that are waiting to execute. A queue can form under a few circumstances. For example, a queue can develop when requests come in for service by the resource at a faster rate than the resource's throughput, or if requests demand longer amounts of time from the resource. A queue can also form if the requests occur at random intervals, such as large batches all at once. When a queue becomes long, work is not handled efficiently, and you might experience delays in response time.

***Response time*** is the measure of time required to do work from start to finish. In a client/server environment you typically measure response time on the client side. Response time generally increases as the load increases. You can measure response time by dividing the queue length for the resource over the resource throughput

# Establishing a Baseline

One of the goals of monitoring Exchange 2000 is to locate problems or anomalies in the server or network. To do this, you must establish a baseline for the various elements you want to monitor. When you collect performance and diagnostic data over a period of time, you can establish a performance baseline. A performance baseline is a range of measurements that represent acceptable performance under typical operating conditions. This baseline provides a reference point that makes it easier to spot problems before they become serious. When you troubleshoot system problems,

performance data gives you information about the behavior of system resources at the time the problem occurs, which is useful in discovering the cause.

A baseline is a measurement that is derived by collecting data over an extended period during varying but typical types of workloads and user connections. The baseline is an indicator of how individual system resources or a group of resources are used during periods of normal activity.

When determining your baseline, it is very important to know the types of work that is done and the days and times when work is done. This helps you associate work with resource usage and to determine whether or not performance during those intervals is acceptable.

For example, if you find that performance diminishes somewhat for a brief period at a given time of day, and you find that at that time many users are logging on or off, it might be an acceptable slowdown. Similarly, if you find that performance is poor every evening at a certain time and you can tell that this time coincides with nightly backups when no users are logged onto the system, then performance loss might be acceptable. But you can make that determination only when you know the degree of performance loss and its cause.

When you build up data on performance over a period of time, with data reflecting periods of low, average, and peak usage, you can make a subjective determination of what acceptable performance is for your system. That determination is your baseline. Use your baseline to detect when bottlenecks develop or to watch for long-term changes in usage patterns that require you to increase capacity.

# Analyzing Results

The baseline you develop establishes the typical counter values you should expect to see when your system performs satisfactorily. However, you need guidelines to help you interpret the counter values and eliminate false or misleading data that might cause you to set your target values inappropriately. You need to identify and investigate bottlenecks to analyze your results and take action.

When you collect and evaluate data to establish a valid performance baseline, you should watch for unusually large values, include ID threads, ignore occasional spikes, use graphs for reporting, exclude start up events, and investigate zero values or missing data.

## Watching for Unusually Large Values

When monitoring processing of the same name, watch for values that are unusually large for one instance and not the other. This can occur because System Monitor sometimes misrepresents data for separate instances of processes of the same name by reporting the combined values of the instances as the value of a single instance. Tracking processes by process identifier can help you solve this problem.

## Including Thread Identifiers

When you are monitoring several threads and one of them stops, the data for one thread might appear to be reported for another. This is because of the way threads are numbered. For example, if you begin monitoring and have three threads, numbered 0,

1, and 2, and if one of them stops, then all remaining threads are sequenced again. This means that the original thread 0 is now gone and the original thread 1 is renamed to 0. As a result, data for the stopped thread 0 can be reported along with data for the running thread 1 because old thread 1 is now old thread 0. To solve this problem, you can include the thread identifiers of the process's threads in your log or display. Use the Thread/Thread ID counter for this purpose.

### Ignoring Occasional Spikes

Do not give too much weight to occasional spikes in data. These might be due to the start up of a process, and they are not an accurate reflection of counter values for that process over time. The effect of spikes can linger over time when using counters that average.

### Using Graphs for Reporting

When monitoring over an extended period of time, use graphs instead of reports or histograms because reports and histograms only show last values and averages. As a result, they might not give an accurate picture of values if you are looking for spikes.

### Excluding Start up Events

Unless you specifically want to include startup events in your baseline, exclude these events because the temporary high values tend to skew overall performance results.

### Investigating Zero Values or Missing Data

Zero values or missing data can impede your ability to establish a meaningful baseline. You should investigate the source of these issues and obtain the missing data, if possible, before you attempt to establish a baseline.

## Identifying and Investigating Potential Bottlenecks

Deviations from your baseline provide the best indicator of performance problems. However, as a secondary reference, Table X.X describes recommended thresholds for object counters. Use this table to help you identify when a performance problem is developing on your system.

| Resource | Object\Counter | Suggested threshold | Comments |
|---|---|---|---|
| Hard disk drive | LogicalDisk\% Free Space | 15 percent | None |
| Hard disk drive | LogicalDisk% Disk Time | 90 percent | None |
| Hard disk drive | PhysicalDisk\Disk Reads/sec, PhysicalDisk\ Disk Writes/sec | Depends on manufacturer's specification | Check the specified transfer rate for your hard disk drives to verify that this rate does not exceed the specifications. In general, Ultra Wide SCSI disks can handle |

| | | | |
|---|---|---|---|
| | | | 50 to 70 I/O operations per second. |
| Hard disk drive | PhysicalDisk\Current Disk Queue Length | Number of spindles plus 2 | This is an instantaneous counter; observe its value over several intervals. For an average over time, use PhysicalDisk\Avg. Disk Queue Length. |
| Memory | Memory\Available Bytes | Less than 4 MB | Research memory usage and add memory if needed. |
| Memory | Memory\Pages/sec | 20 | Research paging activity. |
| Network | Network Segment\% Net Utilization | Depends on type of network | You must determine the threshold based on the type of network you use. For example, for Ethernet networks, 30 percent is the recommended threshold. |
| Paging File | Paging File\% Usage | Above 70 percent | Find the process that is using a high percentage of processor time. Upgrade to a faster processor or install an additional processor. |
| Processor | Processor\Interrupts/sec | Depends on processor, for current CPU's, use a threshold of 1500 interrupts per second. | A dramatic increase in this counter value without a corresponding increase in system activity indicates a hardware problem. Identify the network adapter or hard disk drive controller card causing the interrupts. You might need to install an additional adapter or controller card. |
| Server | Server\Bytes Total/sec | | If the sum of Bytes total/sec for all servers is roughly equal to the maximum transfer rates of your network, you might need to segment the network. |
| Server | Server\Work Item Shortages | 3 | If the value reaches this threshold, consider tuning the InitWorkItems or MaxWorkItems entries in the registry (in HKEY_LOCAL_MACHINE\SYSTEM\ CurrentControlSet\Services\lanmanserver\ Parameters). |
| Server | Server Work Queues\ Queue Length | 4 | If the value reaches this threshold, there might be a processor bottleneck. This is an instantaneous counter; observe its value over several intervals. |
| Multiple Processors | System\Processor Queue Length | 2 | This is an instantaneous counter; observe its value over several intervals. |

Table X.X Recommended Thresholds for Object Counters

**Caution** Do not use a registry editor to edit the registry directly unless you have no alternative. The registry editors bypass the standard safeguards provided by administrative tools. These safeguards prevent you from entering conflicting settings or settings that are likely to degrade performance or damage your system. Editing the registry directly can have serious, unexpected consequences that can prevent the system from starting and require that you reinstall Windows 2000. To configure or customize Windows 2000, use the programs in Control Panel or Microsoft Management Console wherever possible.

Investigating performance problems should always start with monitoring the whole system before looking at individual components. In precise terms, a bottleneck exists if a particular component's limitation keeps the entire system from performing quickly. Therefore, even if one or more components in your system is heavily used, if other components or the system as a whole show no adverse effects, then there is no bottleneck.

For example, suppose that a process has 10 threads, each of which uses exactly 0.999 seconds of processor time once every 10 seconds. If each thread makes a request exactly 1 second after the previous one in perfect sequence, the processor would be 99.9 percent busy, but there is no queue, no interference between the threads, and, technically, no bottleneck, although the system probably can't support any increased load or variation in its request scheduling without creating a bottleneck.

Factors involved in the development of a bottleneck are the number of request for service, the frequency at which requests occur, and the duration of each request. As long as these factors are perfectly synchronized, queues do not develop and bottlenecks do not occur. The device with the smallest throughput is typically the primary source of a bottleneck.

It is difficult to detect multiple bottlenecks in a system. You might spend several days testing and retesting to identify and eliminate a bottleneck, only to find that another bottleneck appears in its place. Only thorough and patient testing of all elements can ensure that you have found all of the problems.

It is not unusual to trace a performance problem to multiple sources. Poor response time on a workstation is most likely the result of memory and processor problems, while servers are more susceptible to hard disk drive and network problems.

Problems in one component might be the result of problems in another component, not the cause. For example, when memory is scarce, the system moves pages of code and data between hard disk drives and physical memory. The memory shortage becomes evident from increased hard disk drive and processor use, but the problem is the lack of memory, not the processor or hard disk drive.